A Closer Look at the Significance of Reliability Specifications

Fred Rybczynski
Storage Technology Corporation
2270 South 88th Street
Louisville, Colorado 80028-4393
email: Fred_Rybczynski@stkamer.storagetek.com
voice: +1-303-673-2976
fax: +1-303-661-7676

ABSTRACT

Device reliability can have a significant impact on the storage and retrieval of computerized data. Poor reliability can cause data to be rewritten multiple times under the control of the device's error detection/correction algorithm. This can result in a reduction in both throughput and capacity. Published reliability claims can also influence the selection of one vendor's product over another. This paper takes a closer look at reliability in order to heighten the reader's awareness of its usage, qualify its significance, and help avoid common misinterpretations.

INTRODUCTION

Deciding which specific component in a system is the most important is somewhat like the very old joke about the parts of the body getting together and arguing as to which is "king." A system, by definition, is an "interdependent group of items forming a unified whole," or stated another way, it's "a set of individual components that form a functional unit." Different components get differing levels of attention, depending on what is going on at the time. For example, you don't pay attention to your nose until it itches or starts running. You don't worry about your eyes until things become blurry. Similarly, you don't worry about your data storage components until something starts going wrong.

Value Added Resellers (VARs) are very sensitive to customer dissatisfaction. They respond quickly to end-user problems because customer satisfaction reflects directly on their company's success. They react rapidly to products that generate too many problem calls and eat into the already-slim profit margin. It is not at all surprising that a recent *Reseller Management* magazine reported that product quality was the number one selection criteria when a VAR decides to carry a vendor's product. Reliability is the single largest contributor to product quality.

Certain measures have been developed to help predict the reliability of data storage products. Some of us have come to expect (and even demand) these predictors. We rely on them for competitive comparisons. We use them to calculate the frequency of preventive maintenance service events and to manage spares parts inventories.

MEASURES OF RELIABILITY

Reliability has been measured and expressed in a number of ways:

- Availability
- Mean Time Between Failure (MTBF)
- Bit Error Rate (BER)
- Failure Rate

Measures of reliability may also have qualifiers associated with them which make them impossible to compare within the category. For example, Bit Error Rate (BER) can be expressed as:

- Corrected bit error rate : media, tape path, electronic data path, ECC algorithm
- Uncorrected bit error rate : theoretical; excludes media contribution
- Undetected bit error rate : theoretical, may include the media contribution

Each unit of measure needs to be defined in a way that is understood and agreed upon by both vendor and customer. Levels of severity also need to be agreed upon. For example, which of the following automobile problems should be considered to be an MTBF "failure"?

- a headlight that no longer works
- a flat tire
- the engine that stops working

DERIVATION OF RELIABILITY SPECIFICATIONS

All reliability numbers start as goals that are extensively influenced by customer desires as well as by the specifications of competitors. An assessment phase should follow the establishment of the goal in order to determine if the goal is attainable and how quickly it can be achieved. Not all vendors go through the assessment phase.

The reliability goals are typically measured and assessed on the basis of

- Reliability of previous products
- Predictive models (MilSpec 217F and Bellcore)
- Reliability data from component suppliers
- Test data of manufactured units
- Field data

During assessment, some vendors may exclude certain data or incidents from reliability computations and/or measurements. An event such as "No Trouble Found" may or may not be included. However, just because a failure is not reproducible does not necessarily mean that a failure did not occur. It's more likely to mean that the vendor has insufficient information to recreate the failure scenario.

What about failures that are attributable to the same unrepaired component? Should that collection of failures be considered a single failure or multiple failures?

It is not sufficient to simply take reliability numbers at face value. Vendors derive them differently, using different qualifiers, resulting in comparisons that are apple-to-orange, not apple-to-apple. Key questions to consider when evaluating reliability claims are:

- How was the number derived ?
- What does it include or exclude ?
- What are the qualifiers ?
- Can it be verified ?
- How much reliability does the application actually require ?

CONTRIBUTORS TO OVERALL RELIABILITY

In the data storage market, the principal contributors to reliability are:

- Duty Cycle
- Operating Environment (operating temperature, temperature cycles, humidity, dirt, vibration)
- Media (quality, storage conditions, age, management)
- Automation
- Redundancy (e.g., RAID)
- Error Recovery
- Design Maturity

CONCLUSION

At its recent "MTBF Shootout" conference, Dataquest pointed out that the data storage market is engaging in an exercise of escalating "specsmanship." MTBF numbers are rapidly spiraling higher. Meanwhile, the device is typically overtaken by the next generation of the technology before the reliability claim can be substantiated empirically.

For example, what does a 1,000,000 hour MTBF really mean? One million hours equates to 114 years at 100% duty cycle. Does anyone actually expect the device to operate non-stop for 114 years? MTBF is intended to predict the rate of device failure, not provide a guarantee of continuous trouble-free performance. The MTBF predicts that 63.2% of deployed devices will fail some time during the specified MTBF period. Unfortunately, MTBF can not predict when the failure will occur. Failure occurs randomly. The higher the MTBF, the more difficult it is to verify empirically because the devices themselves do not remain in use for a long enough time period (e.g., 114 years).

It is important to be aware that differences exist among the various expressions of reliability. Even reliability measures that are labeled identically (e.g., MTBF) can be

calculated differently such that the numbers can not be compared. Ask the questions enumerated previously before comparing reliability numbers from more than one vendor.

STORAGETEK COMMITMENT TO RELIABILITY

StorageTek has a Systems Reliability Group that is charged with the responsibility to assess the performance of our products relative to their initial reliability goals. Their regular assessments of calculated and empirical data are periodically provided to engineering and manufacturing to be factored into continuous process and reliability improvements.